Subband Coding of RF Signals in Reconfigurable Computing Hardware

Christopher M. Brislawn, Scott H. Robinson, and Shane A. Crockett
Los Alamos National Laboratory, Los Alamos, NM 87545–1663 USA

***Abstract*—A subband coding algorithm is designed for compression of wideband RF data and eventual implementation on a reconfigurable computer board. The algorithm involves a multirate filter bank, adaptive uniform scalar quantization, and Huffman entropy coding. We report the performance of the algorithm in software on RF data from the DOE satellites ALEXIS and FORTE. A reconfigurable computer array board, RCA-2, designed and assembled at LANL is presented as a target platform for hardware implementation of the algorithm.**

I. Source Coding for Wideband RF Signals

A. DOE/LANL RF Programs

The U.S. Department of Energy operates a research satellite program in order to fulfill its nuclear weapons treaty verification and nonproliferation responsibilities and to support basic research in the atmospheric and space sciences. There are currently two active satellites in orbit for the DOE program, ALEXIS and the newer FORTE,

`nis – www.lanl.gov/nis – projects/{alexis,forte}.`

Data obtained from the Blackbeard RF receiver on ALEXIS led to the description and analysis in the open scientific literature of the lightning-generated phenomena known as "transionospheric pulse pairs" (TIPPs) [1].

A spectrogram of a TIPP event is shown in Fig. 1. The first (left) impulse is the direct line-of-sight transmission from the point of discharge to the satellite; the second impulse is the ground reflection of the discharge as seen from the satellite. Note that these impulses are *chirped*; in particular, lower frequencies arrive later than higher frequencies. This is a consequence of ionospheric dispersion, which induces a nonlinear, frequency-dependent group delay on the impulses. The nonlinearity is approximated well by the inverse square relationship

$$\tau(f) = c \cdot f^{-2} \quad .$$

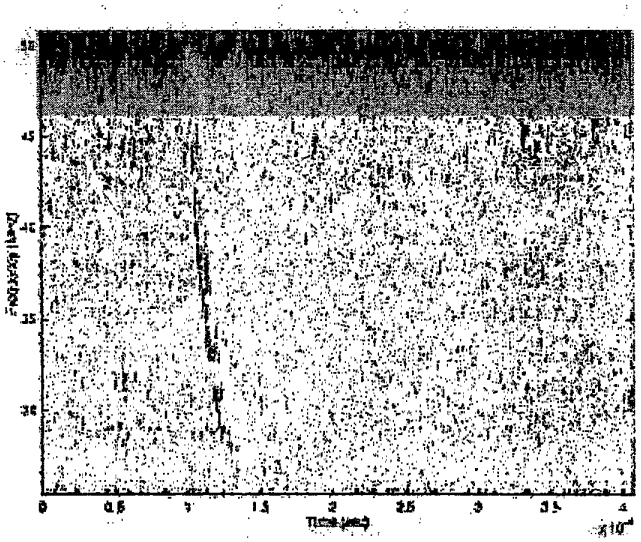

Fig. 1. Spectrogram of a transionospheric pulse pair.

The constant, c, which determines the curvature of the nonlinear chirps, is proportional to the line-of-sight-integrated total electron content (TEC) along the path taken by the impulse through the ionosphere. Accurate TEC estimation is an important objective of scientific data exploitation. Numerical methods for pre-whitening and de-chirping RF spectrograms to estimate TEC and other physical parameters of interest, such as the interpulse lag, are described in [2]. Preserving the accuracy with which TIPP events can be detected and their parameters estimated is the overriding concern of any lossy source coding method applied to these signals.

Consider the bandwidth requirements of the FORTE system. The signal shown in Fig. 1 was acquired by a tunable receiver capable of tuning two independent 22 MHz channels in the 20–300 MHz range. The channels are typically tuned with nominal center frequencies around 38 MHz and 130 MHz. Each channel is digitized to 12 bits at 50 megasamples per second, for a combined output of 150 MB/s. Data must be stored onboard until

Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

e-mail: {brislawn, shr}@lanl.gov; m991368@nadn.navy.mil

0-7803-5491-5/99/$10.00 © 1999 IEEE

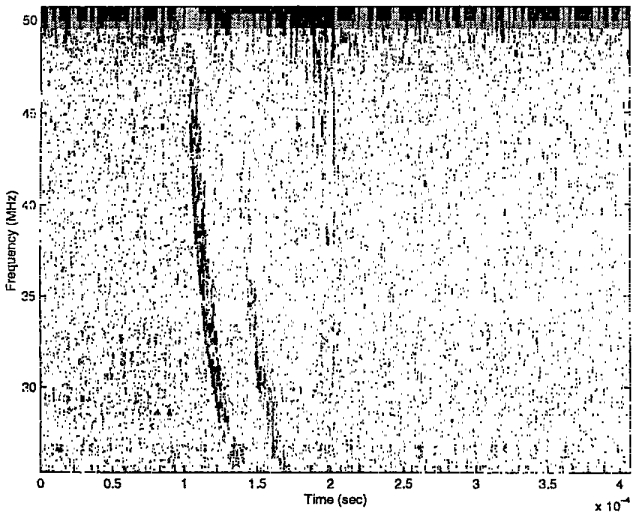

Fig. 2. Spectrogram of TIPP after compression to ~3 bits/sample.

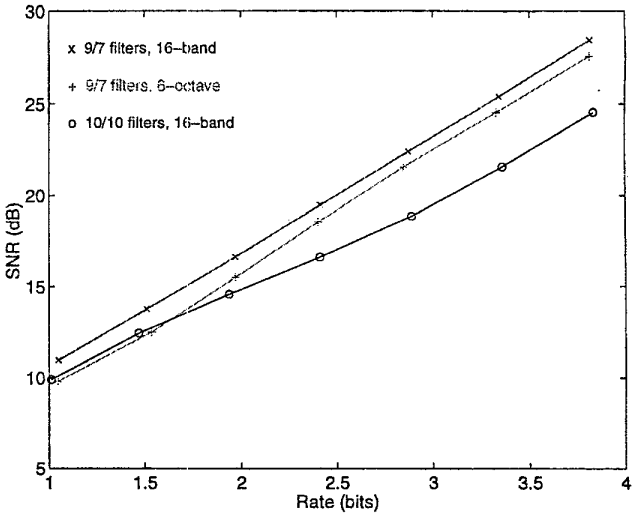

Fig. 3. Rate-distortion performance of different filter bank decompositions.

the satellite is in position to establish a downlink, and since onboard memory is only sufficient for storing ~0.8 seconds of data per channel, recording is triggered by a multichannel event classifier [3].

B. RF Subband Coding Algorithm

The signals described above are being compressed via subband coding schemes. The experiments described here have all been conducted in software. The algorithm does not currently include event detection; the assumption is that we are coding a continuous time-series rather than intermittent bursts of activity. For this reason, the source is chopped into blocks to allow adaptation of the coding to time-varying content. Each block is transformed using a perfect reconstruction multirate filter bank, whose output is buffered for adaptive optimal bit allocation, entropy-constrained scalar quantization, and adaptive Huffman coding. Quantization and entropy coding tables are transmitted as side information with each block. Fig. 2 shows the signal from Fig. 1 after compression to approximately 3 bits/sample (bps) with a block size of 5120 samples, or 102 microseconds (μs) per block.

Optimal bit allocation and scalar quantizer design are accomplished via the method described in Section III of [4]. The bit allocation depends on the variances of the buffered subbands, but we do not employ iterative quantizer design techniques like Lagrangian optimization that require multiple passes through the data. Huffman codebook design is via the single-pass algorithm codified in the JPEG standard [5].

Since the filter bank is non-adaptive, some effort has been expended in trying to determine good transform schemes by testing on training data. As one might expect from the appearance of the spectrogram in Fig. 1, octave-scaled frequency decompositions, or "discrete wavelet transforms" (DWTs), are not necessarily well-suited for this application, in spite of their great success in image coding. This intuition is reinforced by the results shown in Fig. 3, in which we compare the rate-distortion characteristics of a 6-octave DWT decomposition and a uniform-bandwidth 16-band decomposition implemented as a 4-level full binary tree cascade. Both schemes were implemented using the Cohen-Daubechies-Feauveau 9/7 biorthogonal wavelet filter bank [6]. The uniform-bandwidth scheme performed 1/2 to 1 dB better than the 6-octave scheme on RF data. For comparison, another uniform-bandwidth scheme was tested using a related biorthogonal linear-phase wavelet filter bank in which all four filters (both analysis and synthesis) have 10-tap impulse responses. In spite of being somewhat smoother than the wavelets in the 9/7 filter bank, the 10/10 filter bank performed markedly worse on real data. Other filter banks tested against the 9/7 filter bank suffered a similar fate.

To gauge the cost of side information transmitted with each block, tests were conducted on a 4 megasample signal using block sizes ranging from 2048 samples to 65536 samples with the 16-band decomposition and 9/7 filters. (For scale, the 400-μs interval in Fig. 1 contains about 20,000 samples.) The rate-distortion performance of these tests is shown in Fig. 4. As expected, performance drops off as block size decreases and the relative amount of block overhead increases. It is interesting, however, that there appears to be no crossover point at which the greater adaptivity one might have expected from smaller block sizes produced better rate-distortion performance at suitably

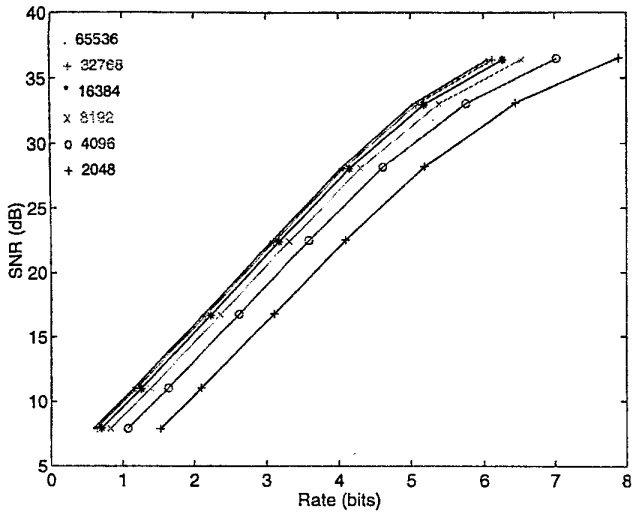

Fig. 4. Rate-distortion performance for various block sizes.

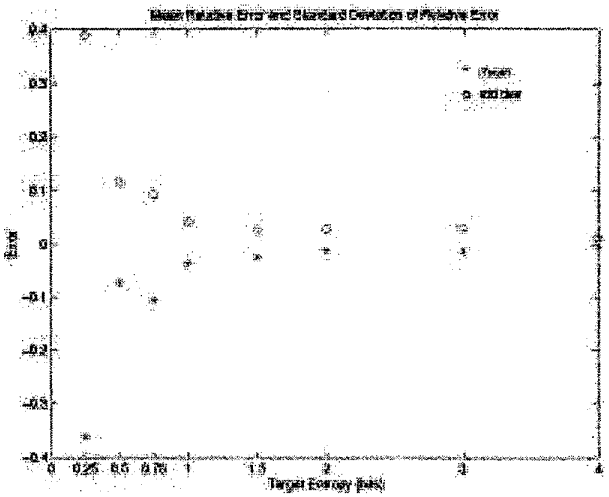

Fig. 5. Mean relative error in estimated TEC for diffuse signals.

high bit rates. At 50 megasamples/sec., a 2048-sample block lasts about 40 μs, which is short enough to localize one of the pulses in Fig. 1 but not short enough to adapt to the time-frequency energy variation in these extremely fast chirps. Moreover, we see in Fig. 4 that the overhead actually *increases* at higher rates due to larger Huffman codebooks. Thus, it appears infeasible to improve the time-frequency adaptivity of the current scheme simply by choosing shorter time blocks unless an alternative can be found that greatly reduces the per-block overhead, such as using adaptive library-based entropy coding. Arithmetic coding is also an attractive option, although its run-time complexity is 2 to 3 times higher than Huffman coding (at least in software).

C. Effects of Quantization on Data Exploitability

The above compression scheme has been subjected to quantitative tests assessing the exploitability of quantized/reconstructed data. Specifically, the parameter estimation tasks described in [2] have been applied to data that has been compressed to rates as low as 1/4 bps. The parameters tested were: estimated TEC, instantaneous dechirped power and SNR, power-autocorrelation $1/e$ width, power-autocorrelation secondary peak SNR, and power-autocorrelation secondary peak lag (i.e., the interpulse lag). Test signals were divided into three categories based on qualitative differences: TIPP signals with two well-localized pulses (e.g., Fig. 1), more diffuse (noisier) TIPP signals, and single-pulse events (no secondary reflection).

While all parameters were reasonably robust with respect to quantization noise at rates down to 3 bps, some parameters, such as power-autocorrelation $1/e$ width, are heavily influenced by loss of signal power due to quantization at lower rates. Other parameters, notably SNR figures, are significantly improved by the high-frequency noise removal that occurs at moderate bit rates. Slant TEC and interpulse lag are very robust with respect to quantization even at rates as low as 0.25 bps.

Fig. 5 plots the relative error in TEC estimated from quantized (versus unquantized, or "true") diffuse signals as a function of target entropy:

$$\text{rel. error} = \frac{\text{TEC}_q - \text{TEC}_t}{\text{TEC}_t}.$$

The values plotted are the mean (*) and standard deviation (o) of the relative error for the "diffuse" category. One can see that TEC is estimated consistently to within a couple percent at rates down to 1 bps. TEC estimation is even more robust for the class of clean TIPP signals: they exhibit less than about 8% relative error even when coded as low as 0.25 bps. Fig. 6 plots the mean relative error in estimated interpulse lag for the category of clean TIPP signals. Lag is estimated with less than 2% error at 0.25 bps.

Complete results of these tests are on the web at:

`www.c3.lanl.gov/wibarf/shane.demo/webfiles`.

II. Reconfigurable Computing

The Space Engineering group at LANL has developed a Reconfigurable Computer Array board (RCA-2). The board was designed to 1) support continuous data throughput of at least 100 MB/s; 2) provide the ability to link multiple boards together into a serial and/or parallel processing structure; and 3) be configured for a wide range of signal and image processing requirements.

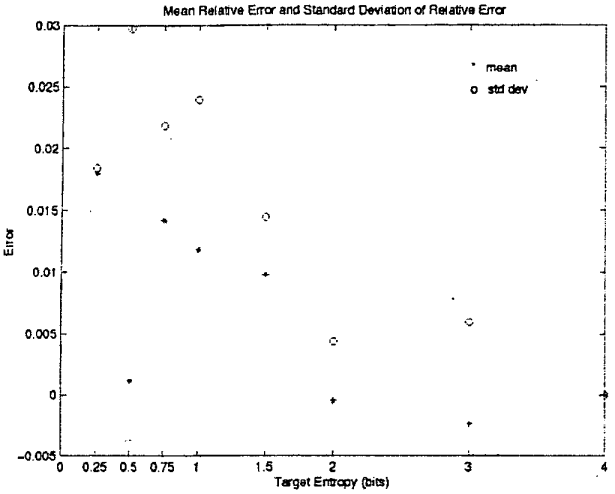

Fig. 6. Mean relative error in estimated interpulse lag for TIPP signals.

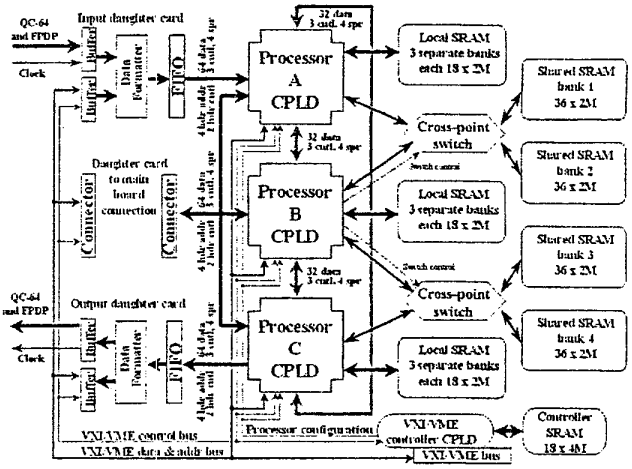

Fig. 7. The RCA-2 reconfigurable computer array board.

To achieve these goals, each reconfigurable processor has a large amount of fast memory to support block oriented processing without resource contentions. The design provides dedicated data paths between processors and from processor to I/O interface to keep blocks of data moving continuously through the board. Finally, the design provides three flexible, high speed I/O ports to allow data to move between boards as required by the processing algorithms. The RCA-2 uses small daughter cards to adapt to different input or output data formats. Fig. 7 shows the RCA-2 board and its major components. The design is implemented on a 'C' size VXI card.

For more details on the RCA-2 design see [7] and

`www.lanl.gov/rcc` .

A. Subband Coding on the RCA-2

Current activity is focused on streamlining the algorithm described in Section I to fit the runtime constraints of the RCA-2. Besides the issue of selecting an effective, implementable entropy coding scheme, a topic of particular concern is the multirate filter bank. To minimize timing problems, it is desirable to run the RCA-2 at clock speeds less than 50 MHz. With 100 megasample/sec. input, this means that we need more than the 2:1 demultiplexing implemented at the front end of a two-channel polyphase realization. This consideration, plus superior attenuation characteristics, are leading us to investigate the use of a cosine-modulated filter bank in place of the 9/7 filter bank cascade. LANL colleague J. Arrowood has demonstrated a 96-tap DFT polyphase filter bank that achieves 60 dB attenuation using a 32-point DFT implementation with 10-bit precision. The same prototype filter in a cosine-modulated filter bank (again with 10-bit precision) acheives about 55 dB attenuation. This is an appropriate range for lossy coding of 12-bit sources. Low-level design and simulation of polyphase and cosine-modulated filter banks for Altera Flex and Xilinx Virtex architectures are now being conducted in concert with John Villasenor's team at UCLA in work supported by DARPA's Adaptive Computing Systems program.

REFERENCES

[1] R. S. Massey and D. N. Holden, "Phenomenology of transionospheric pulse pairs," *Radio Science*, vol. 30, no. 5, pp. 1645–1659, 1995.

[2] A. R. Jacobson, S. O. Knox, R. Franz, and D. C. Enemark, "FORTE observations of lightning radio-frequency signatures: capabilities and basic results," Tech. Rep. LA-UR-98-2046, Los Alamos National Lab, 1998.

[3] K. R. Moore, P. C. Blain, S. D. Briles, and R. G. Jones, "Classification of RF transients in space using digital signal processing and neural network techniques," in *Appl. & Sci. Neural Nets*, vol. 2492 of *Proc. SPIE*, pp. 995–1006, SPIE, 1995.

[4] J. N. Bradley and C. M. Brislawn, "The wavelet/scalar quantization compression standard for digital fingerprint images," in *Proc. Int'l. Symp. Circuits Systems*, vol. 3, (London), pp. 205–208, IEEE Circuits Systems Soc., June 1994.

[5] W. B. Pennebaker and J. L. Mitchell, *JPEG Still Image Data Compression Standard*. New York, NY: Van Nostrand Reinhold, 1992.

[6] A. Cohen, I. C. Daubechies, and J.-C. Feauveau, "Biorthogonal bases of compactly supported wavelets," *Commun. Pure Appl. Math.*, vol. 45, pp. 485–560, 1992.

[7] S. H. Robinson, M. P. Caffrey, and M. E. Dunham, "Reconfigurable computer array: the bridge between high speed sensors and low speed computing," Tech. Rep. LA-UR-98-2510, Los Alamos National Lab, 1998.